

Load management strategy for Particle-In-Cell simulations in high energy particle acceleration

A. Beck^{a,*}, J. T. Frederiksen^b, J. Dérouillat^c

^aLaboratoire Leprince-Ringuet, École Polytechnique, CNRS-IN2P3, Palaiseau 91128, France.

^bNiels Bohr Institute, University of Copenhagen, Blegdamsvej 17, 2100 København Ø, Denmark

^cCEA, Maison de La Simulation 91400 Saclay, France.

Abstract

In the wake of the intense effort made for the experimental CILEX project, numerical simulation campaigns have been carried out in order to finalize the design of the facility and to identify optimal laser and plasma parameters. These simulations bring, of course, important insight into the fundamental physics at play. As a by-product, they also characterize the quality of our theoretical and numerical models. In this paper, we compare the results given by different codes and point out algorithmic limitations both in terms of physical accuracy and computational performances. These limitations are illustrated in the context of electron laser wakefield acceleration (LWFA). The main limitation we identify in state-of-the-art *Particle-In-Cell* (PIC) codes is computational load imbalance. We propose an innovative algorithm to deal with this specific issue as well as milestones towards a modern, accurate high-performance PIC code for high energy particle acceleration.

Keywords: Laser wakefield acceleration, petawatt laser, bubble regime, electron self-injection, relativistic self-focusing, particle-in-cell simulations, Photon-Plasma, SMILEI, dynamic load balancing

1. Introduction

The future CILEX (Centre Interdisciplinaire de la Lumière EXtrême/Interdisciplinary Center for the Extreme Light) is a state-of-the-art laser facility that targets diverse applications using plasmas produced by short (15–30 fs) multi-petawatt (1–10 PW) laser pulses with 0.1–10 Hz repetition rate. It will host the Apollon-10P laser, which will deliver pulses with an instantaneous power up to 10 PW, and the associated infrastructures and experimental setups. It will thus offer an opportunity for scientific breakthrough in various domains [2]. The facility is located in France (Paris area) and is expected to open to the international user community in 2016.

The experimental setups and the required instruments for relevant diagnostics remain to be designed. New regions of the parameter space are now accessible with Apollon-10P. They have never been reached before and large numerical simulation campaigns are necessary to explore them. These campaigns have three objectives: (1) to make sure that the most promising laser and plasma parameters are accessible in CILEX, (2) to improve the preparation in terms of instrumentation and radioprotection, and (3) to help understanding the physics at work.

Computational support is therefore mandatory for such large experimental facilities, and PIC simulations are the best candidates for this task. A large variety of PIC codes exists. They range from simple serial codes making severe physical assumptions to speed up computations, to elaborate and computation-

ally costly massively parallel PIC codes, supposedly more accurate.

In Section 2, we present a first numerical study of a typical CILEX case of LWFA using three standard but very different codes, to understand their strengths and limitations. The LWFA is at the core of the CILEX scientific program, with the principle task being to explore the experimental prospects of using ultra-short PW pulses for multi-GeV acceleration; this will occur in the bubble regime [3], taking advantage of electron self-injection and laser pulse self-guiding in a low-pressure gas/plasma cell. Section 3 discusses computational load imbalance, how it heavily slows down LWFA simulations but also how this barrier can be lifted with an appropriate dynamic load balancing algorithm.

2. Self injection in the bubble-out regime at CILEX

2.1. The codes

This section details the results given by the PIC codes Wake, Calder-Circ (CC) and Photon-Plasma (PP) for a standard CILEX LWFA experiment.

Wake is a quasi-static, axi-symmetric code [17] with test particle tracking [10–13]. In spite of its restrictive context, Wake can be used to identify the nonlinear optical processes at work in the laser-plasma interaction and have been used, for instance, to explain the details of pulse evolution for acceleration beyond the theoretical pulse depletion limit[1]. Wake is a sequential code and a typical LWFA simulation only takes half a day of computation on a single core.

*Corresponding author.

Email address: beck@llr.in2p3.fr (A. Beck)

CC is an explicit quasi-3D PIC code [15]. It uses poloidal mode decomposition for the electromagnetic fields and currents, while computing super-particles (SP) trajectories in the 3D Cartesian space. Using the two lowest-order modes is sufficient to accurately capture the most important aspects of the interaction [15]. CC has been used in many laser-plasma interaction studies where long laser propagation (a few cms) were needed[1, 4, 10]. Indeed, the mode decomposition allows to get 3D results for the cost of a couple of 2D standard simulations and is well adapted to long propagations. A typical LWFA simulation with CC takes several days of computation on 500 cores.

PP is an explicit fully 3D PIC code. It benefits from high-order schemes and has demonstrated excellent scalability up to 250,000+ cores using a hybrid MPI+OpenMP approach[8]. It has been used mostly in astrophysical cases[5] and was recently upgraded to support long laser propagations. For this application, its high-order scheme is especially valuable. With a 6th order scheme, one can conserve stability and suppress numerical dispersion (for a given wavelength) by simply setting the appropriate resolution in time and space. In contrast, numerical dispersion in a traditional 2nd order scheme cannot be reduced without improving temporal and spatial resolution, which is orders of magnitude more costly. Another benefit of the sixth order scheme is that the numerical dispersion, when not suppressed, tends to overestimate the group velocity of the laser and therefore limits the numerical Cherenkov effect with respect to lower order schemes who tends to underestimate the laser group velocity.

2.2. Simulation parameters

The simulation presented here uses the expected laser parameters of the first CILEX shots with the so called “F2” beam. The plasma parameters have been optimized as shown in [1] and with a complementary parametric study with Wake. They provide a good trade off between electron energy, energy spread and total charge of the electron bunch. The carrier wavelength of the Apollon-10P laser is $\lambda_0 = 0.8 \mu\text{m}$. The laser pulse is Gaussian in all directions has a duration of $\tau_L = 25 \text{ fs}$ (FWHM in intensity), the laser strength parameter is $a_0 = 5.55$ and its transverse spot size is $w_0 = 16.4 \mu\text{m}$. The pulse energy on target is 7.41 J for a peak power of 0.28 PW. And finally, the plasma density is set to $n_e = 1.4 \times 10^{18} \text{ cm}^{-3}$ starting by a 0.3 mm long ramp and the laser is focused at the foot of the ramp.

Similar numerical parameters were chosen for PP and CC for benchmarking purposes. The longitudinal and transverse resolutions in both codes are respectively set to $\Delta z = 0.125 [c/\omega_0]$ and $\Delta x, y = 1.5 [c/\omega_0]$, where ω_0 is the laser frequency. The time steps Δt differ for stability reasons and are respectively 0.0625 and 0.124 $[1/\omega_0]$ for PP and CC. Wake uses the coarser resolution $\Delta z = 0.47 [c/\omega_0]$, $\Delta x = 5.55 [c/\omega_0]$ and $\Delta t = 15.42 [1/\omega_0]$.

2.3. Results

The simulations were run long enough to give a quantitative evaluation of the entire injection process at play. The most important results are summarized in Figure 1. After the entry of




Figure 1: Evolution of the amplitude of the driver, injected charge and bubble radius given by the three different codes. On the bottom panel, plain lines show the evolution of a_0 and dashed lines the injected charge. Wake does not give any information about the injected charge because it assumes a quasi-static evolution.

the laser pulse in the plasma density ramp, the laser self-focuses and the so-called bubble structure forms in its wake. The bubble quickly shrinks to its minimum size before starting to grow again very slowly as injection begins. The three codes show a similar scenario, with, in particular, a similar relativistic self-focusing of the pulse. However, even though the driver is well described in all codes, the bubble dynamics and the total charge of injected electrons differ. After 1400 μm of propagation, 275 pC are injected in PP whereas only 200 pC are injected in CC. This 30% difference is pretty small with respect to the variations observed experimentally from shot to shot. It could be explained by the fact that CC is based on a cylindrical geometry and the longitudinal axis is subject to user-defined boundary conditions which might not always be accurate and could artificially influence the injection. It is also known that the numerical Cherenkov effect introduces a significant amount of noise on axis in CC[14]. Nevertheless, even though here CC gives results very similar to PP, full 3D simulations would still be mandatory in cases where the cylindrical symmetry is broken such as laser pulses with experimental aberrations, the presence of an external magnetic field or a transverse density gradient.

3. The importance of computational load-management

Fully 3D PIC simulations are rarely used and only by groups with massive access to top-tier supercomputers. Their high cost is often prohibitive. In some cases, the raw computational cost can be significantly reduced by the use of a Lorentz boosted frame[16]. But even in these favorable cases, they still require a very large number of cells and SP and fit only on massively




Figure 2: Total number of super-particles owned by each MPI process in a full 3D Photon-Plasma simulation of LWFA. It ran on 2048 nodes of the Fermi system at Cineca. Note the base-10 logarithmic scaling. The initial value is $\text{Log}_{10}(N_{init,MPI}) = 5.19$ — or, equivalently, $N_{init,MPI} = 155\,520$. The minimum and maximum are $N_{min,MPI} = 2570$, and $N_{max,MPI} = 6\,070\,158$, respectively.

parallel systems. It is almost inevitable that, at some point of the run, a small part of the system becomes overloaded with too many SP, slowing down the entire simulation. At the same time, a large part of the system is underloaded and remains idle, wasting precious computational resources. This phenomenon, known as load imbalance, is more likely to happen as the size of the system and the number of spatial dimensions increases. Massively parallel 3D cases being most problematic as seen in Figure 2. This very large simulation is a remarkable illustration of the limitation of actual 3D LWFA simulations. Some areas of the simulation are almost completely empty of SP whereas some are very densely populated and handle 40 times as many SP as they did in the initially uniform plasma. A slow-down by a similar factor is therefore expected, and may lead to unacceptable under-utilization of the computational resources.

The use of OpenMP is often believed to be a solution to load imbalance in PIC codes. More OpenMP threads implies fewer and larger MPI domains and consequently smoother SP distribution between them. Since the load is directly related to the number of SP, it is also shared more homogeneously. Nevertheless, in the case of Figure 2, 16 OpenMP threads per MPI process were used and yet a very strong imbalance remains. The OpenMP effectiveness is limited on large systems because as the number of cores increases, the number of OpenMP threads remains the same, bounded by the number of hardware threads per CPU. As a consequence, the number of MPI processes must increase and the relative sizes of the MPI domains shrink, eventually becoming smaller than the load variation scale. At this point, load imbalance hits and performances are lost. OpenMP can delay this situation, but not prevent it. Consequently, imple-




Figure 3: Example of a 32×32 patches decomposition between 7 MPI processes. MPI domains are delimited by different uniform colors. The line shows the Hilbert curve and the dots the center of the patches. The curve goes through all patches. It starts from the patch with coordinates (0, 0) and ends at patch with coordinates (31, 0).

mentation of some form of dynamic load balancing to achieve scalability, is paramount. Below, two innovative approaches are considered: dynamic load balancing via “patching” and dynamic load limiting via k-means clustering. These methods are not mutually exclusive, rather, they are complementary.

3.1. Dynamic load balancing

The “patching” method was first described in [7] and has been implemented in SMILEI [18], a PIC code developed to support the CILEX community. It consists in breaking the large structures encapsulating fields and SP data of each MPI process into a multitude of smaller independent structures called “patches”. The benefit is that the code now has a data structure particularly well adapted to thread parallelization and, in particular, to OpenMP.

Patches are coherently organized along a 1D Hilbert space filling curve[9]. Each MPI process owns a segment of this curve. Its properties of continuity and locality ensure that the obtained domains are compact and minimize synchronization communications. This decomposition has a very flexible organization and can be very different from a Cartesian decomposition as shown in Figure 3. Taking advantage of this flexibility, MPI processes are able to dynamically exchange patches in order to even out their respective loads.

Results for a simple 2D LFWA simulation with SMILEI running on 24 nodes of the OCCIGEN system are compared in Figure 4. As shown in [7], load imbalance occurs early in the simulation. In the LWFA case, this sets in already at bubble formation around iteration 6000. The first few hundreds iterations take around 0.5 s each, independent of parallelization strategy.

In a run with a pure MPI parallelization, the time needed for 10 iterations increases up to 80 s which gives an imbalance ratio of approximately 14. In the hybrid run with 12 OpenMP




Figure 4: Evolution of the time spent for 10 SMILEI cycles along a LWFA run on 24 OCCIGEN nodes. “MPI” is a pure MPI run, 576 MPI processes. “Hybrid” is a hybrid MPI+OpenMP run, 48 MPI processes + 12 OpenMP threads per process. “Hybrid balanced” is the same hybrid run with the dynamic load balancing active.

threads, this ratio is approximately halved and is only around 2 when the dynamic load balancing is activated.

3.2. Dynamic load limiting

The other approach is based on computational SP splitting/merging in low/high SP number density cells, respectively. It has been implemented in PP. The merging/splitting must carefully take into consideration conservation laws of physics; the scheme is costly, but precise. It is based on an accelerated k-means clustering method. Besides being of importance in load imbalanced situations, merging and splitting of SP is of profound significance when carrying out simulations in which detailed Monte Carlo particle-particle collisions, ionization or quantum electrodynamics (QED) are implemented. Otherwise, they can lead to an excessive number of SP and to memory bound situations because of the tremendous amount of new SP created; smaller SP of the same species from collisions, electrons in the case of ionization and electron-positron pairs or even photons from QED. In this regard, this approach can be considered more as a dynamic load limiting algorithm rather than balancing. It is absolutely mandatory for modern simulations since these new kind of developments are emerging in both laboratory and astrophysical plasmas communities. Details of this approach can be found in [6].

4. Summary and outlook

The importance of 3D simulations has been established. Their cost is still prohibitive for most applications since large scale simulations are so massively parallel that load imbalance becomes a barrier. It cannot be alleviated by brute force parallelism. That is merely a waste of computational resources. Germaschewski et al. suggested a new data structure for PIC

codes, adapted to both modern massively parallel hardware and to the associated algorithmic challenges [7]. This structure has been adapted with success in SMILEI [18] and shows tremendous improvement in the case of LWFA, even in 2D.

Recent advances in the physics included in PIC codes in both astrophysical plasmas and laser-plasma communities often leads to prohibitive numbers of SP. Now, the problem is not just only compute bound, but also memory bound. This issue may be resolved only through a load-limiting algorithm such as particle merging through k-means clustering [6].

The two methods above could be combined into a coherent dynamic load management strategy without interfering with any of the physics included in the code. It is even compatible with different geometry models like CC itself which is also subject to load imbalance.

Acknowledgments

PP and CC simulations were performed using the PRACE Research Infrastructures resources FERMI based in Italy at CINECA and Curie based in France at TGCC under the grant 2013091859. SMILEI simulations were performed using DARI allocation c2015067484 on the OCCIGEN system, based in France at CINES. This work was partially supported with a bilateral STSM visit grant under the EU COST Action MP 1208 “Developing the physics and the scientific community for inertial fusion” and by the LabEx PALM (ANR-10-LABX-0039-PALM) under project SIMPLE.

References

- [1] A. Beck, Y. Kalmykov, X. Davoine et al., Physical processes at work in sub-30 fs, PW laser pulse-driven plasma accelerators: Towards GeV electron acceleration experiments at CILEX facility, Nucl. Instrum. Methods in Phys. Res. A (2013), 10.1016/j.nima.2013.11.003.
- [2] B. Cros, B. S. Paradkar, X. Davoine et al., Laser plasma acceleration of electrons with multi-PW laser beams in the frame of CILEX, Nucl. Instrum. Methods in Phys. Res. A (2013), 10.1016/j.nima.2013.11.003.
- [3] A. Pukhov, J. Meyer-ter-Vehn, Appl. Phys. B: Lasers Opt. 74 (2002) 355.
- [4] B. M. Cowan, S. Y. Kalmykov, A. Beck et al., J. Plasma Phys. 78 (2012) 469–482.
- [5] J. T. Frederiksen, Astrophys. J. Lett. 680 Issue 1 (2008).
- [6] J. T. Frederiksen, G. Lapenta, M. E. Pessah, Particle control in phase space by global k-means clustering, 2015, arXiv:1504.03849.
- [7] K. Germaschewski, W. Fox, N. Ahmadi et al., The plasma Simulation Code: A modern particle-in-cell code with load-balancing and GPU support, 2013, arXiv:1310.7866.
- [8] T. Haugbølle, J. Trier Frederiksen, Å. Nordlund, Phys. Plasmas 20 (2013) 062904.
- [9] D. Hilbert, Math. Ann. 38, 459 (1891).
- [10] S. Y. Kalmykov, S. A. Yi, A. Beck et al., New J. Phys. 12 (2010) 045019.
- [11] S. Y. Kalmykov, S. A. Yi, A. Beck et al., Plasma Phys. Control Fusion 53 (2011) 014006.
- [12] S. Y. Kalmykov, A. Beck, S. A. Yi et al., Phys. Plasmas 18 (2011) 056704.
- [13] S. Y. Kalmykov, A. Beck, X. Davoine, E. Lefebvre, and B. A. Shadwick, New J. Phys. 14 (2012) 033025.
- [14] R. Lehe et al., Phys. Rev. STAB 16 (2013) 021301.
- [15] A. F. Lifschitz, X. Davoine, E. Lefebvre et al., J. Comput. Phys. 228 (2009) 1803.
- [16] J.-L. Vay, Phys. Rev. Lett. 98 (2007) 130405.
- [17] P. Mora and T. M. Antonsen, Jr., Phys. Plasmas 4 (1997) 217.
- [18] Smilei web page <http://www.maisondelasimulation.fr/projects/Smilei/html/index.html>.